


# Preserving Data Integrity Assurance While Providing High Levels of Compression for Primary Storage

December, 2007



# Table of Contents

<b>Abstract</b> .....	<b>2</b>
<b>At-a-Glance</b> .....	<b>2</b>
The Technology .....	2
Complete Data Integrity .....	2
<b>Protecting Data Integrity While Improving Storage Capacity</b> .....	<b>2</b>
Safeguarding Your Data .....	2
Storage Capacity and Data Integrity – No Trade-Off .....	2
Data Integrity Concerns When Integrating a Storage Capacity Solution .....	2
Using the Storwize STN Compression Appliance in a NAS Environment .....	2
<b>How the STN Appliance Ensures Data Integrity</b> .....	<b>3</b>
Preserving Existing Data Integrity Mechanisms in the Network .....	3
Ensuring Reliable Compression and Decompression .....	3
Compressing Data Using the LZ Compression Algorithm .....	5
Reliability at the Hardware Level – ECC memory modules .....	5
<b>Testing Data Integrity Mechanisms</b> .....	<b>5</b>
<b>Advantages of the Storwize STN Appliance</b> .....	<b>5</b>
<b>About Storwize</b> .....	<b>6</b>

## Abstract

Storwize has created a solution that provides high levels of data compression for primary storage. In most cases this compression is achieved with a performance gain making the Storwize solution viable for even high performance transaction processing environments. As part of the design of the solution Storwize has developed the DIA Data Integrity Assurance architecture. This document describes the DIA Data Integrity Assurance architecture and shows how the DIA Data Integrity Assurance is able to ensure end to end data integrity.

The multi tiered data protection provided by the DIA Data Integrity Assurance and built into every Storwize data compression appliance is a quantum leap forward in guaranteeing data integrity.

## At-a-Glance

### The Technology

- DIA Data Integrity Assurance architecture
- Multi-core architecture
- Industry-proven compression algorithms
- Average compression of 65% across all data types
- Primary storage compression solution
- Real-time random access to compressed data
- Application agnostic solution

### Complete Data Integrity

To ensure 100% data integrity, the STN product line:

- Preserves existing data integrity mechanisms in the storage array
- Preserves existing data integrity mechanisms in the network
- Verifies the data can be decompressed before writing it to storage
- Verifies the reliability of storing and retrieving the compressed data
- Uses the industry-proven Lempel-Ziv (LZ) lossless data compression algorithm
- Uses error-correcting code (ECC) memory modules to ensure data integrity for memory access

## Protecting Data Integrity While Providing High Levels of Data Compression

### Safeguarding your data

In any IP network environment, components and systems can fail at any time, potentially causing costly disruptions to your organization's IT infrastructure. In designing a network, administrators ensure that reliable data integrity mechanisms are in place to protect against damaging network outages, software failures, virus threats and even human error.

When your storage capacity needs grow and it becomes necessary to expand your IP network, only a storage solution that provides proven reliability without interfering with these data integrity mechanisms is acceptable.

### Storage Capacity and Data Integrity – No Trade-Off

A storage network offers a balance between performance and availability, based on the specific requirements of the organization. Any storage capacity optimization solution that compresses primary data must integrate seamlessly into the data integrity mechanisms already operating in the network.

The Storwize solution accomplishes all of this, protecting your data while increasing your storage capacity.

As an integral part of the design of the STN solution Storwize developed the DIA Data Integrity Assurance architecture that ensures that Storwize products provide end-to-end verification throughout the data path, fault detection and correction and verification of decompression reliability.

The multi-tiered data protection provided by DIA Data Integrity Assurance and built into every Storwize data compression appliance is a quantum leap forward in guaranteeing data integrity.

### Data Integrity Concerns When Integrating a Storage Capacity Solution

A component in the data path, such as a compression appliance integrated into a storage network infrastructure, must address the following data integrity concerns:

- **Verification of decompression reliability** – A component that compresses primary data must verify the reliability of the data decompression process before acknowledging the write to the client. Data that an application writes must return to the application in exactly the same form when it is read.
- **End-to-end verification** – Write acknowledgements from the storage system must be passed to the client only after the data is stored on the storage array.
- **Fault detection and correction** – Any network component must preserve the multiple levels of data integrity assurances inherent in the network.

### Using the Storwize STN Compression Appliance in a NAS Environment

A principal advantage of the network-attached storage (NAS) environment is the reliability it provides in a complex data-sharing environment. This paper describes how the STN compression appliance integrates seamlessly into NAS environments, increasing the storage capacity, while assuring data integrity by:

- Preserving the data integrity mechanisms already in place in the network
- Providing mechanisms to guarantee the reliability of the compressed data through the DIA Data Integrity Assurance architecture

## How the STN Appliance Ensures Data Integrity

When you install the STN appliance, all of the data integrity assurance mechanisms inherent to the network continue to operate as usual, with the STN appliance performing its role transparently. The STN appliance verifies that each data block that it compresses can be restored to the original form. The following sections describe how the STN appliance:

- Ensures reliable compression and decompression using the DIA Data Integrity Assurance architecture
- Preserves the existing data integrity mechanisms in the network
- Uses the reliable LZ compression algorithm
- Uses ECC memory modules to ensure memory read/write reliability

### Preserving Existing Data Integrity Mechanisms in the Network

The STN appliance preserves data integrity mechanisms at the protocol and storage levels.

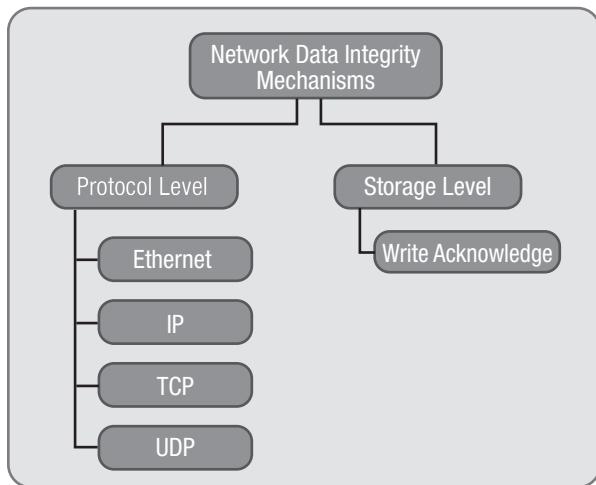


Figure 1. Network Data Integrity Mechanisms

### Ensuring Data Integrity at the Network Protocol Level

Protocol-level integrity checks ensure data integrity during transmission between components in your network. These checks are cyclic redundancy check (CRC) checksum mechanisms that verify data integrity during transmission of data between components.

When you install the STN appliance, the checksum mechanisms employed in your network continue to function as usual. The STN appliance manages two sets of checksum mechanisms:

- Checksums on the uncompressed data in the connection between the STN appliance and the switch
- Checksums on the compressed data in the connection between the STN appliance and the storage array

The network protocol-level checksums include:

- **Ethernet-layer checksum (IEEE 802.3)** – Protects data against corruption during transmission on the Ethernet wire by performing a CRC checksum at the frame level
- **IP header checksum (RFC 791)** – Ensures the integrity of the data fields in the IP header
- **TCP checksum (RFC 793)** – Verifies the integrity of each data block during TCP/IP transmission
- **UDP checksum (RFC 1122)** – Optional checksum that verifies the integrity of data packets associated with the UDP datagram

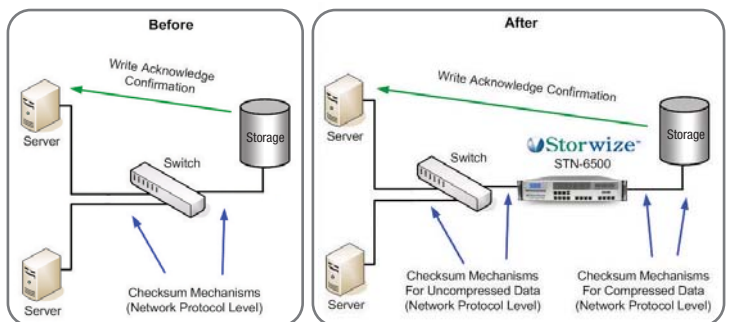


Figure 2. Preserving Data Integrity Mechanisms in the Network

### Ensuring Data Integrity at the Storage Level

The write acknowledge is a critical aspect of the functionality of client server storage applications. When the storage system completes a client write request, the storage system returns a write acknowledge message. The write acknowledge message provides confirmation that the data was transmitted successfully through the network and written to storage. This confirmation protects against failures in transmission caused by any component of the network, between the client server and the storage system.

When the STN appliance is installed in your network, it transparently passes write acknowledge messages received from the storage system to the client that requested the write, ensuring that the data was stored on the storage array.

### Ensuring Reliable Compression and Decompression

After selecting the most reliable compression algorithm to use, the next step is to execute it flawlessly. Storwize has developed a unique DIA Data Integrity Assurance architecture for the STN appliance product line to verify the accuracy of every data block that the appliance compresses. The DIA Data Integrity Assurance architecture:

- Verifies that the compressed data can be decompressed with absolute reliability
- Verifies the integrity of the compressed data during a read operation

The following steps describe the DIA Data Integrity Assurance architecture process.

### Step 1: Compression

When the STN appliance receives a raw data block to write, it applies the compression algorithm to create a compressed data block. During this compression step, the STN appliance stores a copy of the raw data block in memory.




Figure 3. DIA Data Integrity Assurance Step 1 – Compressing the Data

### Step 2: Decompression

After creating the compressed data block, the STN appliance immediately decompresses it to restore the data to its original form. The decompressed version of the data is compared bit-by-bit with the copy of the original raw data stored in memory to ensure that the two versions are identical. When the process is complete, the STN appliance has verified that the compressed data can be decompressed with proven reliability.

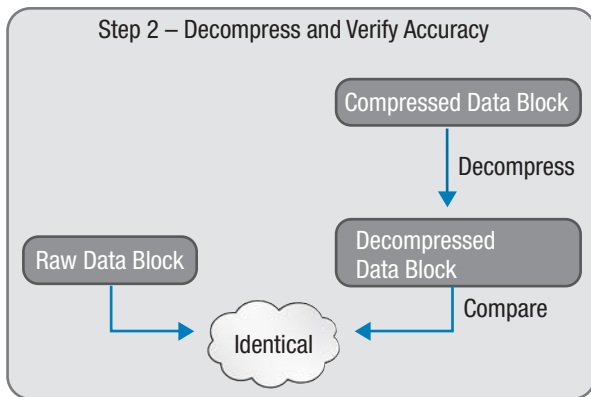


Figure 4. DIA Data Integrity Assurance Step 2 – Decompressing and Verifying the Data

### Step 3: Checksum

After verifying that compressed data can be decompressed reliably, the STN appliance calculates a checksum for the raw data block. It stores this checksum with the compressed data on the storage system. The checksum is used for verifying the integrity of the stored data when it is read back to the STN appliance during a read operation.

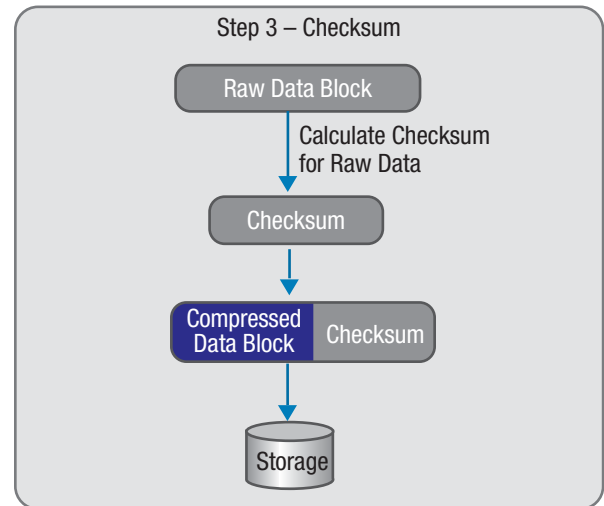


Figure 5. DIA Data Integrity Assurance Step 3 – Checksum Stored with the Compressed Data

### Verifying Storage Data Integrity Using the DIA Data Integrity Assurance

The checksum stored with the compressed data block provides an important integrity mechanism, which is utilized during the read operation. When a server requests data, the compressed data is retrieved from the storage system and sent to the STN appliance. The STN appliance decompresses the data and recalculates the checksum for the decompressed data. It then confirms that the new checksum is identical to the original value stored with the data. This mechanism verifies that the data remained error-free during all of the following operations:

- **Write** – Transmission of the compressed data to the storage system during the write operation
- **Storage** – Storage of the compressed data on the storage system
- **Read** – Transmission of the compressed data from the storage system to the STN appliance during a read operation

After confirming that the data is error-free, the STN appliance sends the data in its original form to the server that sent the request for the data.

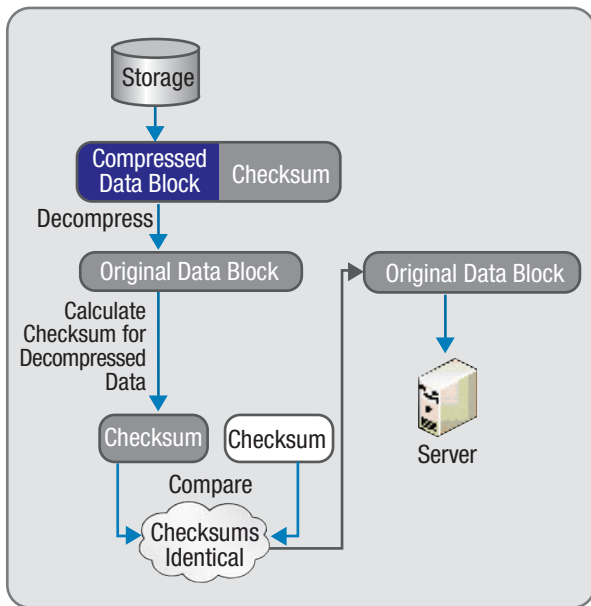


Figure 6. Verification of Storage Integrity by Checksum

### Compressing Data Using the LZ Compression Algorithm

The STN appliance employs the industry-proven Lempel-Ziv (LZ) lossless data compression algorithm. The LZ algorithm is known for its absolute reliability and its compression effectiveness.

The LZ algorithm has a long track record in the industry, serving as the basis of numerous compression algorithm variations. The algorithm has proven itself in a wide range of software and hardware products, including the familiar zip file format and GNU zip (gzip). The LZ algorithm is also used in the archival storage tape systems that organizations rely on for their last-resort data backup. In this critical role, the algorithm has proven itself to be absolutely reliable.

Using the proven LZ compression algorithm is an important part of assuring data integrity during the compression and decompression of data.

### Reliability at the Hardware Level – ECC memory modules

To ensure data integrity at the hardware level while the data passes through the STN appliance's memory modules, the appliance utilizes parity error-correcting code (ECC) memory. This memory is the most reliable type of memory modules available. It provides a built-in checksum method to ensure the accuracy of data being written to and read from memory.

### Testing Data Integrity Mechanisms

Storwize tests the STN appliance rigorously in a variety of network environments to ensure that the appliance you incorporate into the network retains total reliability of data integrity mechanisms. As part of this testing, Storwize simulates errors caused by network and environmental factors, to test the appliance under the most rigorous conditions.

## Advantages of the Storwize STN Appliance

Installing the STN appliance in a storage configuration provides unmatched benefits to the storage environment:

- Primary storage compression – The STN appliance compresses the data before it reaches the storage unit. This provides unique advantages over other storage capacity solutions:
  - Improves response time – Reduces traffic to and from the storage system
  - Reduces the storage system's read/write activity – Compressing data before it reaches the storage system enables the storage system to store and provide data with less read and write activity.
  - Read gains – Increases the effective size of the storage read cache, improving the performance of read operations. Storage system read performance is greatly influenced by the cache hit ratio. Better use of the read cache reduces the need for disk head movement, which is slower than cache access. Because Storwize compresses the data before it reaches the storage system, the effective size of the read cache increases. The read cache improvement is proportionate to the compression ratio.
  - Write gains – Increases the effective size of the write cache, improving the performance of write operations. Storage systems typically cache write operations to enable immediate acknowledgement to clients and reduce response time. The size of the write cache significantly impacts the throughput of write operations to storage. The STN appliance compresses the data before it reaches the write cache, so the effective size of the cache increases proportionately with the compression ratio.
- Reduces infrastructure costs – Using the STN appliance reduces the expensive infrastructure requirements of expanding your storage capacity, including:
  - Floor space – savings of up to 95%
  - Electricity savings of up to 90%
  - Cooling savings of up to 75%
- No changes or server resources required – The STN appliance provides storage gains with the advantages of self-contained, appliance simplicity:
  - No need to change the storage system configuration
  - No need to change your workflow
  - No server resources used

## About Storwize

Storwize, Inc. is the only provider of real-time data compression solutions that dramatically boost storage space in network storage environments. Storwize appliances are installed between the organization's storage arrays and their network. Designed to work with any storage array it is connected to, the Storwize appliance typically provides organizations with 2-5 X additional capacity, with associated savings in capital, operational, maintenance, and management resources. The company recently received funding from Sequoia Capital, one of the largest and most influential venture capital firms. Storwize solutions have been shipped to customer sites worldwide and successfully tested for interoperability with leading storage vendors, such as EMC and NetApp. Storwize has won several coveted industry honors from analyst firms and leading IT publications, and has added many new customers in a variety of vertical markets worldwide. Storwize is a member of the Green Grid. For more information, visit the company website: [www.storwize.com](http://www.storwize.com).



**Storwize Inc.**, 2033 Gateway Place, 5th Floor, San Jose, CA 95110, USA | Toll free: (800) 670-0281 | Tel: (408) 961-7373 | Fax: (408) 961-7375

**E-mail:** [info@storwize.com](mailto:info@storwize.com)

**[www.storwize.com](http://www.storwize.com)**

© 2007 Storwize, Inc. all rights reserved. The information published herein is provided "as is" and is subject to change without prior notice. Storwize makes no representations or warranties of any kind with respect to the information and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. The information disclosed herewith and any derivatives thereof are confidential and solely owned by Storwize, Inc. The disclosure or usage of any of the information disclosed without the affirmative prior written consent of Storwize, Inc. for any purpose and to any third party is strictly restricted. Storwize is a trademark of Storwize, Inc.. All other trademarks used herein are the property of their respective owners. All trademarks and trade names, including but not limited to "Storwize", "STN-5000", "STN-6000", "PrediSave", "DIA Data Integrity Assurance", either defined or identified, marked as <sup>TM</sup> or unmarked are and will remain at all times, the sole property of Storwize, Inc.